

REMARKSI. Introduction

In response to the Office Action dated July 31, 2003, no claims have been cancelled, amended or added. Claims 1, 3-11, 13-21 and 23-30 remain in the application. Re-examination and re-consideration of the application is requested.

II. Prior Art RejectionsA. The Office Action Rejections

In section (2) of the Office Action, claims 1, 3-7, 11, 13-17, 21 and 23-27 were rejected under 35 U.S.C. §103(a) as being obvious over Schiefer et al., U.S. Patent No. 5,761,653 (Schiefer) in view of Chiang, U.S. Patent No. 6,477,523 (Chiang). In section (3) of the Office Action, claims 6-10, 16-20, and 26-30 were rejected under 35 U.S.C. §103(a) as being unpatentable over Schiefer in view of Chiang, U.S. Patent No. 6,477,523 (Chiang) and further in view of Raitto et al., U.S. Patent No. 5,991,754 (Raitto).

Applicant's attorney respectfully traverses these rejections.

B. The Applicant's Claimed Invention

Applicant's claimed invention, as recited in independent claims 1, 11, and 21, is generally directed to a method of optimizing execution of a query that accesses data stored on a data store connected to a computer. Claim 1 is representative and recites the steps of generating cardinality estimates for one or more query execution plans for the query using statistics of one or more automatic summary tables that vertically overlap the query, and using the generated cardinality estimates to determine an optimal query execution plan for the query.

C. The Schiefer Reference

Schiefer describes a method for estimating cardinalities for query processing in a relational database management system. The method is suitable for use with a query optimizer for estimating cardinalities for sets of columns or keys resulting from a grouping operation or a duplicate removal operation.

D. The Chiang Reference

Chiang describes a method, apparatus, and article of manufacture for generating statistics for use by a relational database management system. A global aggregate spool is generated for each of a plurality of partitions of a subject table that are spread across a plurality of processing units of a computer system. Each of the global aggregate spools is scanned to generate summary records. The summary records are then merged to generate interval records for a compressed histogram of the subject table, wherein the compressed histogram includes both equal-height intervals and high-biased intervals. The compressed histogram can then be analyzed to estimate the cardinality associated with one or more search conditions of a user query or other SQL statement. Compared to a strictly equal-height histogram, the compressed histogram allows the relational database management system to more accurately estimate the cardinality associated with various search conditions. As a result, the relational database management system can better optimize the execution of the user query.

E. The Raitto Reference

Raitto describes a method and system for processing queries, where the queries do not reference a particular materialized view. Specifically, techniques are provided for handling a query that specifies a first set of one or more aggregate functions, where the particular materialized view reflects a second set of one or more aggregate functions. Whether the query can be rewritten is determined based on the aggregate functions in the first and second sets, and the corresponding arguments. Techniques are also provided for processing a query that (1) does not reference a particular materialized view, (2) specifies a first set of one or more aggregate functions, where the particular materialized view reflects a second set of one or more aggregate functions. A technique is also provided for rewriting queries that specify an outer join that has a dimension table on the child-side of the outer join and a fact table on the parent-side of the outer join. The query is rewritten to produce a rewritten query by replacing references to the fact table in the query with references to a materialized view. The rewritten query specifies an outer join that has the dimension table on the child side and the materialized view on the parent side.

F. Applicant's Claimed Invention Is Patentable Over The Cited References

Applicant's claimed invention is patentable over Schiefer, Chiang and Raitto, because it includes a combination of limitations not taught or suggested by the cited references, taken individually or in any combination.

The combination of Schiefer and Chiang is cited by the Office Action as teaching all of the steps or elements of the independent claims 1, 11 and 21.

Applicant's attorney disagrees.

The Office Action states that Chiang teaches the elements "generating cardinality estimates for one or more query execution plans for the query using statistics of one or more automatic summary tables that vertically overlap the query" at col. 6, lines 32-40 and 50-65. However, at the indicated locations, Chiang merely states the following:

Chiang: Col. 6, lines 32-40

According to the preferred embodiment of the present invention, a new kind of database statistics, known as a compressed histogram, are generated for use by the Optimizer subsystem of the PE 114 in optimizing an execution plan. The compressed histogram includes high-biased intervals and/or equal-height intervals that allow the Optimizer subsystem of the PE 114 to more accurately estimate the cardinality associated with various conditions of the execution plan.

Chiang: Col. 6, lines 50-65

A first collection step is responsible for building a global aggregate spool and a sequence of summary records on each AMP 116 participating in the statistics collection (i.e., on each AMP 116 that manages a partition of the subject table), wherein multiple copies of the first collection step are executed simultaneously and in parallel by the AMPs 116. In this manner, the global aggregate spool may be considered partitioned in the same manner as the subject table.

Each row of the global aggregate spool includes: (1) a distinct value from the partition of the subject table and (2) the number of rows in the partition of the subject table having the distinct value. The global aggregate spool is considered global in the sense that a distinct value from the subject table can only be found on a single AMP 116, because the subject table is partitioned across multiple AMPs 116.

Nothing in the above description from Chiang can fairly be said to represent "generating cardinality estimates for one or more query execution plans for the query using statistics of one or more automatic summary tables that vertically overlap the query."

In Chiang, summary records are constructed from a global aggregate spool. Each row of the global aggregate spool includes: (1) a distinct value from the partition of the subject table and (2) the number of rows in the partition of the subject table having the distinct value. Each summary record

includes: (1) a sort key, (2) a distinct value, and (3) the number of rows in the partition of the subject table having the distinct value.

However, the summary records in Chiang are not "automatic summary tables" or "materialized views." As noted in Applicant's specification, automatic summary tables are pre-computed queries.

Also, Chiang does not determine that an automatic summary table vertically overlaps a query. An automatic summary table vertically overlaps a query when the set of predicates applied by the automatic summary table is a subset of the predicates required by the query.

However, there is no discussion of vertically overlapping automatic summary tables in Chiang. Indeed, Chiang is directed only to the construction of a compressed histogram of a subject table without reference to a query.

Consequently, Chiang does not teach or suggest "generating cardinality estimates for one or more query execution plans for the query using statistics of one or more automatic summary tables that vertically overlap the query."

The Office Action also states that Schiefer teaches the elements "using the cardinality estimates to determine an optimal query execution plan for the query" at col. 3, lines 37-60. However, at the indicated location, Schiefer merely states the following:

Schiefer: Col. 3, lines 37-60

It is another object of the present invention to produce a better cardinality estimate by utilizing information and attributes which can be obtained from the catalog for the relational database management system. The additional information includes cardinalities for existing unique keys, column equivalence classes, functional dependencies, statistical functional dependencies, and statistically unique keys.

In a first aspect, the present invention provides a method for estimating cardinalities for a key formed from a grouping of columns in a table for use in a query optimizer for a relational database management system, wherein selectivities and keys associated with columns in the table are provided in a catalog, said method comprising the steps of: (a) determining an equivalence class for each column in said key; (b) for each said equivalence class determining an effective cardinality for each of said columns belonging to said equivalence class; (c) determining a cardinality for each of said equivalence classes by choosing the minimum effective cardinality for the columns belonging to said equivalence class; and (d) estimating a cardinality value for said key from the product of said cardinalities for said equivalence classes.

Nothing in the above description from Schiefer can fairly be said to represent "using the cardinality estimates to determine an optimal query execution plan for the query."

In the context of Applicant's claims, the cardinality estimates are generated using statistics of one or more automatic summary tables that vertically overlap the query.

In Schiefer, however, the cardinality estimates are generated by (1) determining an equivalence class for each column in a key; (b) for each equivalence class, determining an effective cardinality for each of the columns belonging to the equivalence class; (c) determining a cardinality for each of the equivalence classes by choosing the minimum effective cardinality for the columns belonging to the equivalence class; and (d) estimating a cardinality value for the key from the product of the cardinalities for the equivalence classes.

Consequently, Schiefer does not teach or suggest "using the cardinality estimates to determine an optimal query execution plan for the query."

Raitto does not overcome the deficiencies of Schiefer and Chiang. Recall that Raitto was cited only against dependent claims 6-10, 16-20 and 26-30, and is specifically directed to queries that do not reference a particular materialized view (automatic summary table).

Consequently, even when combined, the Schiefer, Chiang and Raitto references teach away from Applicant's invention. Moreover, the various elements of Applicant's claimed invention together provide operational advantages over the cited references. In addition, Applicant's invention solves problems not recognized by the cited references.

Thus, Applicant submits that independent claims 1, 11 and 21 are allowable over Schiefer, Chiang and Raitto. Further, dependent claims 3-10, 13-20 and 23-30 are submitted to be allowable over Schiefer, Chiang and Raitto in the same manner, because they are dependent on independent claims 1, 11 and 21, respectively, and because they contain all the limitations of the independent claims. In addition, dependent claims 3-10, 13-20 and 23-30 recite additional novel elements not shown by Schiefer, Chiang and Raitto.

### III. Conclusion

In view of the above, it is submitted that this application is now in good order for allowance and such allowance is respectfully solicited.

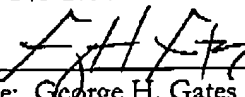
Should the Examiner believe minor matters still remain that can be resolved in a telephone interview, the Examiner is urged to call Applicant's undersigned attorney.

Respectfully submitted,

GATES & COOPER LLP  
Attorneys for Applicants

Howard Hughes Center  
6701 Center Drive West, Suite 1050  
Los Angeles, California 90045  
(310) 641-8797

Date: October 31, 2003

By:   
Name: George H. Gates  
Reg. No.: 33,500

GHG/sjm

G&C 30571.272-US-U1